

Claims

What is claimed is:

1. A method for administering the workload of a database system as it executes one or more requests, the method including:

5        sorting the one or more requests into one or more workload groups, each workload group having an associated level of service desired from the database system;

executing the one or more requests in an order intended to achieve the levels of service associated with each of the workload groups;

assigning system resources to the one or more workload groups as necessary to provide the 10      level of service associated with each workload group;

monitoring on a short-term basis the execution of requests to detect a deviation from the level of service greater than a short-term threshold and if such a deviation is detected:

adjusting the assignment of system resources to workload groups to reduce the deviation;

monitoring on a long-term basis to detect deviations from the expected level of service greater 15      than a long-term threshold, and if such a deviation is detected:

adjusting the execution of requests to better provide the expected level of service.

2. The method of claim 1 where adjusting the assignment of system resources to applications includes:

adjusting the CPU and associated IO allocation assigned to each workload group.

- 20     3. The method of claim 1 where sorting includes:

dividing the requests into one or more workload groups; and

assigning service level goals (SLGs) to the one or more workload groups.

4. The method of claim 3 where sorting further includes:

mapping the one or more workload groups to the one or more classes depending on the SLGs 25      assigned to each of the one or more workload groups.

5. The method of claim 3 where assigning SLGs includes:

accepting input from a user; and

providing guidance to the user.

6. The method of claim 5 where providing guidance includes:
  - receiving information regarding the performance of the system; and
  - providing guidance to the user based on the received information regarding a current ability of the system to satisfy the SLGs.
- 5 7. The method of claim 3 where assigning SLGs includes:
  - publishing the SLGs to the system.
8. The method of claim 1 where monitoring on a short-term basis includes:
  - monitoring the throughput of requests assigned to each workload group; and
  - calculating a performance goal index (PGI) from throughput for each workload group.
- 10 9. The method of claim 8 where monitoring the throughput includes measuring the average response time, where an SLG includes a response time goal, and where calculating the PGI includes:
  - dividing the average response time by the response time goal.
10. The method of claim 1 where adjusting the assignment of system resources includes:
  - adjusting the assignment of the system resources to the one or more workload groups to minimize the maximum PGI.
- 15 11. The method of claim 1 where the workload groups are divided into high priority workload groups and low priority workload groups and where adjusting the assignment of system resources includes:
  - adjusting the assignment of the system resources in favor of the high priority workload groups to minimize the maximum PGI for the high priority workload groups.
- 20 12. The method of claim 1 where adjusting the execution of requests includes:
  - swapping out a request based on its workload group assignment to free up system resources.
13. The method of claim 1 where adjusting the execution of requests includes:
  - aborting the execution of a request based on its workload group assignment.
- 25 14. The method of claim 1 where adjusting the execution of requests includes:
  - delaying execution of a request based on its workload group assignment.

15. The method of claim 1 where monitoring on a long-term basis includes:  
logging the deviations greater than the long-term threshold in an error log.
  16. The method of claim 15 where monitoring on a long-term basis further includes:  
making the error log available for a user to view.
- 5    17. A computer program, stored on a tangible storage medium, for use in administering the workload of a database system as it executes one or more requests, the program including executable instructions that cause a computer to:
- sort the one or more requests into one or more workload groups, each workload group having an associated level of service desired from the database system;
  - 10    execute the one or more requests in an order intended to achieve the levels of service associated with each of the workload groups;
  - assign system resources to the one or more workload groups as necessary to provide the level of service associated with each workload group;
  - monitor on a short-term basis the execution of requests to detect a deviation from the level of service greater than an short-term threshold and if such a deviation is detected:
  - 15    adjust the assignment of system resources to workload groups to reduce the deviation;
  - monitor on a long-term basis to detect deviations from the expected level of service greater than a long-term threshold, and if such a deviation is detected:
  - adjust the execution of requests to better provide the expected level of service.
- 20    18. The computer program of claim 17 where when adjusting the assignment of system resources to workload groups the executable instructions cause the computer to:
- adjust the CPU and associated IO allocation assigned to each workload group.
19. The computer program of claim 17 where when sorting the executable instructions cause the computer to:
- 25    divide the requests into one or more workload groups; and  
assign service level goals (SLGs) to the one or more workload groups.

20. The computer program of claim 19 where when sorting the executable instructions further cause the computer to:

map the one or more workload groups to the one or more classes depending on the SLGs assigned to each of the one or more workload groups.

5 21. The computer program of claim 19 where when assigning SLGs the executable instructions cause the computer to:

accept input from a user; and  
provide guidance to the user.

10 22. The computer program of claim 21 where when providing guidance the executable instructions cause the computer to:

receive information regarding the performance of the system; and  
provide guidance to the user based on the received information regarding a current ability of the system to satisfy the SLGs.

15 23. The computer program of claim 20 where when assigning SLGs the executable instructions cause the computer to:

publish the SLGs to the system.

20 24. The computer program of claim 17 where when monitoring on a short-term basis the executable instructions cause the computer to:

monitor the throughput of requests assigned to each workload group; and  
calculate a performance goal index (PGI) from throughput for each workload group.

25. The computer program of claim 24 where monitoring the throughput includes measuring the average response time, where an SLG includes a response time goal, and where when calculating the PGI the executable instructions cause the computer to:

divide the average response time by the response time goal.

26. The computer program of claim 17 where when adjusting the assignment of system resources the executable instructions cause the computer to:

adjust the assignment of the system resources to the one or more workload groups to minimize the maximum PGI.

27. The computer program of claim 17 where the workload groups are divided into high priority workload groups and low priority workload groups and where when adjusting the assignment of system resources the executable instructions cause the computer to:

adjust the assignment of the system resources in favor of the high priority workload groups to  
5 minimize the maximum PGI for the high priority workload groups.

28. The computer program of claim 17 where when adjusting the execution of requests the executable instructions cause the computer to:

swap out a request based on its workload group assignment to free up system resources.

29. The computer program of claim 17 where when adjusting the execution of requests the executable instructions cause the computer to:

10 abort the execution of a request based on its workload group assignment.

30. The computer program of claim 17 where when adjusting the execution of requests the executable instructions cause the computer to:

delay execution of a request based on its workload group assignment.

15 31. The computer program of claim 17 where when monitoring on a long-term basis the executable instructions cause the computer to:

log the deviations greater than the long-term threshold in an error log.

32. The computer program of claim 31 where when monitoring on a long-term basis the executable instructions further cause the computer to:

20 make the error log available for a user to view.

33. A database system including:
- a massively parallel processing system including:
- one or more nodes;
- a plurality of CPUs, each of the one or more nodes providing access to one or more CPUs;
- a plurality of data storage facilities each of the one or more CPUs providing access to one or more data storage facilities;
- a process for execution on the massively parallel processing system for administering the workload of a database system as it executes one or more requests, the process including:
- sorting the one or more requests into one or more workload groups, each workload group having an associated level of service desired from the database system;
- executing the one or more requests in an order intended to achieve the levels of service associated with each of the workload groups;
- assigning system resources to the one or more workload groups as necessary to provide the level of service associated with each workload group;
- monitoring on a short-term basis the execution of requests to detect a deviation from the level of service greater than a short-term threshold and if such a deviation is detected:
- adjusting the assignment of system resources to workload groups to reduce the deviation;
- monitoring on a long-term basis to detect deviations from the expected level of service greater than a long-term threshold, and if such a deviation is detected:
- adjusting the execution of requests to better provide the expected level of service.
34. The database system of claim 33 where adjusting the assignment of system resources to workload groups includes:
- adjusting the CPU and associated IO allocation assigned to each workload group.
35. The database system of claim 33 where sorting includes:
- dividing the requests into one or more workload groups; and
- assigning service level goals (SLGs) to the one or more workload groups.

36. The database system of claim 35 where sorting further includes:
  - mapping the one or more workload groups to the one or more classes depending on the SLGs assigned to each of the one or more workload groups.
37. The database system of claim 35 where assigning SLGs includes:
  - 5 accepting input from a user; and
  - providing guidance to the user.
38. The database system of claim 37 where providing guidance includes:
  - receiving information regarding the performance of the system; and
  - providing guidance to the user based on the received information regarding a current ability of
- 10 the system to satisfy the SLGs.
39. The database system of claim 35 where assigning SLGs includes:
  - publishing the SLGs to the system.
40. The database system of claim 33 where monitoring on a short-term basis includes:
  - monitoring the throughput of requests assigned to each workload group; and
- 15 calculating a performance goal index (PGI) from throughput for each workload group.
41. The database system of claim 40 where monitoring the throughput includes measuring the average response time, where an SLG includes a response time goal, and where calculating the PGI includes:
  - dividing the average response time by the response time goal.
- 20 42. The database system of claim 33 where adjusting the assignment of system resources includes:
  - adjusting the assignment of the system resources to the one or more workload groups to minimize the maximum PGI.

43. The database system of claim 33 where the workload groups are divided into high priority workload groups and low priority workload groups and where adjusting the assignment of system resources includes:

adjusting the assignment of the system resources in favor of the high priority workload groups  
5 to minimize the maximum PGI for the high priority workload groups.

44. The database system of claim 33 where adjusting the execution of requests includes:

swapping out a request based on its workload group assignment to free up system resources.

45. The database system of claim 33 where adjusting the execution of requests includes:

aborting the execution of a request based on its workload group assignment.

10 46. The database system of claim 33 where adjusting the execution of requests includes:

delaying execution of a request based on its workload group assignment.

47. The database system of claim 33 where monitoring on a long-term basis includes:

logging the deviations greater than the long-term threshold in an error log.

48. The database system of claim 47 where monitoring on a long-term basis further includes:

15 making the error log available for a user to view.